

Application Form

Institution Name

Sheffield Hallam University

Originating Department

Learning and IT Services

Contact Name (and email address)

Dave Thornley (d.h.thornley@shu.ac.uk)

Objective of the VMware ESX service

The objective of the VMware server virtualisation service is to improve the provision of servers used to provide the university's IT services in the following areas:

- Cost effective use of hardware
- Disaster Recovery/Business Continuity capability
- Infrastructure protection and management
- Flexibility and agility in service delivery

The service came about following a review of university datacentres in late 2003. The review highlighted a number of issues with the model for server deployment in use at the time, particularly

- A high number of small, underutilised servers typically running a single application
- A rising cost of deploying servers as gigabit ethernet and SAN based storage were increasingly used
- Rising electricity costs
- A looming capacity issue with floor and rack space, power and cooling capacities.

The model was attractive however as it allowed a high degree of isolation of applications from each other, making support and management of services simple.

Following the review two possibilities were considered: server consolidation onto fewer larger servers or a shift in server provision to the use of virtual machines for production services.

It was felt that the consolidation option would result in server configurations that were highly complex, as many services would be deployed onto each server. Diagnosing and resolving incidents and problems in applications could potentially have a high impact on other services running on the same hardware. This would potentially force routine support work out of hours, or result in service interruptions during required hours.

Virtualisation was felt to be the preferred option as it would continue to allow a high degree of application isolation, retain the simple configuration of servers and still address the issues identified in the data centre review. Both VMware

Workstation and VMware GSX Server (now VMware Server) were already used for production services in the university and the recently released VMware ESX Server 2.0 was identified as a potential solution.

A paper outlining the issues and the potential risks and benefits of ESX was submitted to the University Executive Group (Appendix 1). The proposal was accepted, and a pilot of VMware ESX 2.0 deployed to evaluate it as a platform for production services. Following a review after 12 months (Appendix 2), the service was expanded and moved to production use.

Description of the VMware ESX service

VMware ESX is now the first choice platform for new services. Servers are deployed onto physical machines only if they cannot be run as virtual machines. Reasons for this include

- high hardware requirements. It is of questionable value to virtualise a highly utilised resource.
- reliance on the service by VMware. This avoids circular dependencies and prevents problems managing the service should problems develop.
- lack of support by vendors for use on VMware. This needs to be considered for critical services.

Three years on the university has 170 virtual machines, representing just over half the university's server estate. Over two thirds of the virtual machines are running production applications. VMware has delivered all the objectives identified and provided some significant benefits to the university:

- Major savings in the development and deployment of new services
- Reduced environmental footprint of running the datacentre, both power consumption and equipment purchased and subsequently disposed of
- Improved hardware utilisation, with CPU usage at 40% or higher
- Significant improvements in the disaster recovery capability of the university's IT services
- Improved support for innovation, development and testing of services

These benefits are described in more detail below

Cost effective use of hardware

Previously the university had a large number of servers with average CPU utilisation below 5%. These servers cost on average £5,000 to buy and £1,000 per year to run (hardware and software maintenance, power and cooling). Using VMware, each virtual server costs £3,000, a 40% improvement, and an overall saving to date of £350,000. Annual virtual server running costs are around £50 per server per year, an improvement of 95%. As more virtual machines are deployed the cost effectiveness of the service improves as economies of scale are exploited.

An additional significant benefit of VMware is the reduction of the environmental impact of the datacentres. A physical server typically uses 1.2 amps of electricity, a VMware host around 2 amps due to the higher utilisation. Over a year of 24x7 operation 170 physical servers would use

686,000KWh, the 9 VMware ESX hosts only 60,500KWh. In a year this saves over £43,000 and 269 tons of CO₂. The university also requires less infrastructure (switches, racks etc) to support its operation so uses and disposes of less equipment.

Disaster Recovery and Business Continuity

The university has two datacentres on separate sites and services are designed to failover between them where possible. VMware hosts are installed on both sites and use SAN storage mirrored between sites in real time.

In a traditional server deployment services are deployed with redundant capacity to provide disaster recovery and business continuity. This is a highly effective but expensive solution resulting in underutilisation of hardware. Using the university's virtual infrastructure virtual machines can be started and run on any host. The use of mirrored disk ensures that any virtual machine can be available on either site as required. As VMware host capacity would be reduced following a major incident non-production virtual machines would be shut down to free up capacity to run production services. Prioritising resources on hosts will ensure that essential services can operate normally if capacity is reduced. Any server on VMware acquires this disaster recovery provision without needing redundant servers. This model has been tested several times when electrical work in the university has required several complete shutdowns of datacentres. Services based on VMware have been able to run normally throughout the shutdowns.

Infrastructure protection and management

The implementation of effective change management for IT services has been a major initiative in the university for some time. It requires representative test environments to verify changes to service configurations behave as expected and to enable the development of new services. Previous test environments were based on old equipment often retired from the datacentres. As a result they were not representative of the hardware in use or the live configuration of services and required significant time to configure before a valid test could be carried out. On VMware test environments can be prepared as required by cloning live virtual servers, generating an exact copy of the hardware and configuration to test changes.

As virtual machines are available free to users and can be created within hours of a request VMware enables new services to be developed and tested with little risk to the university both financially, as no new hardware is needed and in terms of availability of services as testing and development can be completely isolated from the production environment.

Flexibility and agility in service delivery

The introduction of VMware has dramatically increased the flexibility and responsiveness of Learning and IT Services. Virtual machines are provided free to users and can be ready just hours after a request is made. Without VMware it would typically take several weeks and require budgetary authority and a financial process. This has enabled a significant improvement in

support for innovation and development of new services for the university, and fast deployment of servers whenever required. In several occasions virtual machines have been quickly deployed to replace failed hardware as it was a faster option than waiting for an engineer to arrive.

VMware allows the prioritisation of server resources and so makes it easier for the university to change the priority of services over time. Services required in the working day can be allocated additional resources during core hours, then the priority changed to focus resources on services supporting remote working during evenings and weekends. If services are being phased out the level of resource they are allocated can be reduced over time as usage drops. When they are no longer needed they can be 'turned off' and archived on disk in case needed again. With physical servers this would require keeping old or redundant servers commissioned in case needed again when a service was removed.

Supporting documentation about VMware ESX

[Appendix 1](#) – Proposal for implementation of virtualisation of production services

[Appendix 2](#) – Review of the VMware ESX pilot

Appendices [3i](#) and [3ii](#) – Vendor case studies on VMware ESX at Sheffield Hallam University

Name of Staff involved (including job titles and email addresses)

Dave Thornley

Service Support Manager, IT Services

d.h.thornley@shu.ac.uk

Nik Sellars

Platform Support Team Leader, IT Services

n.sellars@shu.ac.uk

Chris Hailey

Distributed Systems Team Leader, IT Services

c.r.hailey@shu.ac.uk

Richard Smith

Senior Analyst/Programmer, IT Services

r.smith@shu.ac.uk

Sam Duckfield

Analyst/Programmer, IT Services

s.d.duckfield@shu.ac.uk

Alan Court

Analyst/Programmer, IT Services

a.j.court@shu.ac.uk

Support of Institution UCISA Representative

Name: John Hemingway